

FEATURE ARTICLE

Tom Napier

Considering an Analog Filter

As the ever-changing tech market continues to top itself, Tom reminds us not to forget the classics. Right now, all eyes are on digital processing, but analog filters are still useful. After hearing about the various types and benefits, you'll find analog filters simple to boot.

n

ot all signal filtering problems can be solved with a DSP chip and some fancy software. An analog filter often does a better job and needn't be hard to design.

In an electronics world dominated by digital processing it's easy to forget that analog filters still have a part to play. Digital filters are compact, their performance is closely specified, and they can implement filter types that simply don't exist in the analog domain. Unfortunately, they process numbers not signals.

THE LIMITS OF DIGITAL

Before you can use a digital filter, you have to convert real-world signals into numerical form. Analog-to-digital conversion has a limited resolution and dynamic range. It also introduces quantization noise. The filtered result often must be converted back to analog form before it can do anything useful. But, D/A conversion introduces more errors and generates a spurious high-frequency output; this could defeat the original object of the filter.

The whole process—sampling, processing, and reconversion—depends on clock frequency. Not only does this limit the highest input frequency that

can be handled, it also introduces aliasing, a phenomenon with no equivalent in an analog filter. Aliasing arises because sampling the input inherently converts signals and noise from frequencies higher than the sampling rate into the frequency band of interest. Unless frequencies higher than half of the sampling rate are filtered out of the input signal before conversion, they appear distorted in the output.

The reverse process occurs during the output conversion. Wanted signals are aliased into the frequency region higher than the output clock rate. Thus, a practical digital filter may require analog filters both at its input and output to do its job. That's when you have to ask yourself whether or not it would be simpler and less expensive to stay in the analog domain.

WHY FILTER?

A filter passes the band of frequencies in which your signal of interest lies and rejects noise and interfering signals. For example, radio receivers use narrowband intermediate frequency filters to pick out just one signal. Filters also can shape signal waveforms to adapt them for further processing. Space data receivers use pulse-shaping filters to maximize the output signal-to-noise ratio.

Filters can be built to pass only low frequencies, only high frequencies, or a band of frequencies. You can build a notch filter that passes all frequencies except a particular one. Filters can be passive (i.e., containing only inductors, capacitors, and resistors) or they can be active, using amplifiers and feedback connections.

CLASSIC FILTERS

Analog filters used to be designed as if they were matched transmission lines. Many inductors and capacitors were connected together with a termi-

	Butterworth	Bessel
L1	0.50650	0.26439
L2	1.56543	1.05449
C1	0.63880	0.39739
C2	1.97432	1.71653

Table 1—Component values for a four-pole LCR filter are normalized to -3 dB at 1 rad/s.

nation resistor at one end. The more inductors and capacitors you added, the sharper you could make your filter discriminate between wanted and unwanted signals. Filter design became complicated. Most engineers simply looked up pre-calculated designs in a huge book that gave the component values and performance of each configuration.

Inductance/capacitance (LC) filters for audio frequencies were bulky and expensive. The early equipment to frequency-multiplex many phone lines on one cable needed ferrite-cored inductors up to 2" in diameter. Then, in the 1950s, two engineers, R. P. Sallen and E. L. Key, discovered how to make low-pass filters from resistors, capacitors, and unity-gain (cathode-follower) amplifiers. The introduction of these active filters greatly decreased the size of telephone exchange equipment.

Active filters are easier to design, build, and tune than traditional LC filters. Because they incorporate buffer amplifiers, sections can be linked in series without mutual interaction. The basic Sallen and Key filter shown in Figure 1 is a two-pole, low-pass stage, an ideal building block for many filter designs. Figure 2 shows how it behaves as the damping is changed. To get predictable results, the buffer amplifier must have a bandwidth many times greater than the filter.

Active filters have two limitations, they can't easily be made to work above about 50 MHz and they can be overloaded by interfering signals. If there's a 1000x stronger interfering signal on a frequency near your wanted signal, a passive filter could separate the two. In an active filter, the strong signal would saturate the amplifier, destroying the smaller signal.

FILTER TYPES

Early filter designers were working with audio modulation of radio frequency carriers, so generally they wanted to pass a narrow range of frequencies while rejecting all others. Their ideal was the brick wall filter and much ingenuity went into

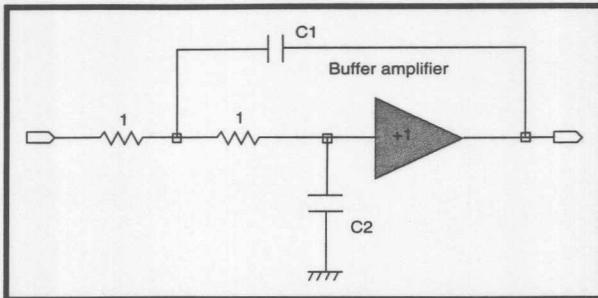


Figure 1—The basic Sallen and Key filter has two poles. This filter can be designed for a variety of corner frequencies and damping factors and is easily assembled into multi-pole filters.

approximating it. These days we more often transmit broadband digital signals. Preserving the signal shape is more important than sharply chopping off high frequencies.

This brings us to Murphy's Law of filters: A filter with sharp frequency cutoff causes ringing, which distorts the shape of the signal. To get clean rising and falling edges, you must use a constant delay (a.k.a., linear phase) filter. This has a rounded frequency response.

There is a compromise. You can build a sharp cutoff filter and follow it with an all-pass filter, which has no effect on the amplitude of a signal but does phase shift and delay it. A sharp cutoff filter has a low delay at low frequencies but a high delay near the cutoff frequency. You can add an all-pass filter with a high delay at low frequencies to equalize the delay. Equalization

can easily double the complexity of a filter so it's not done unless you absolutely must have both a sharp cutoff and constant delay. (Uniform delay makes digital filters popular.)

The basic filter is a low-pass filter that removes high frequencies from the signal. If you want a high- or band-pass filter, you usually start by designing a low-pass filter, then transform it mathematically to give the configuration and component values for the filter you want. Over the years, people invented filters having different characteristics, each optimized for a different task. These classic filters are known by the names of their inventors—Butterworth, Bessel, Paynter, and Chebychev. The exception is the elliptical filter.

The Butterworth filter is the basic low-pass filter. It has a flat response up to its cutoff frequency and then rolls off smoothly. But, it distorts the signal edges by introducing ringing. If signal shape is important, the best filter is the Bessel filter. This filter has a constant delay for all frequencies below its cutoff frequency. Although it delays and rounds a signal, it doesn't change its fundamental shape. Unfortunately, achieving this delay characteristic forces it to roll off slow-

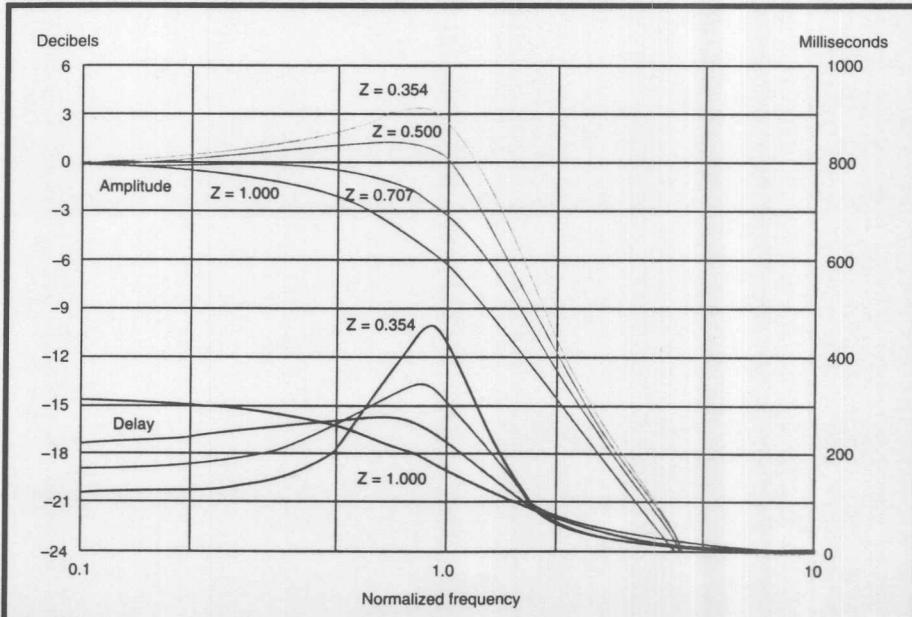


Figure 2—The amplitude response and delay of a Sallen and Key filter vary with the damping factor. The product of the two capacitors sets the corner frequency and their ratio controls the damping.

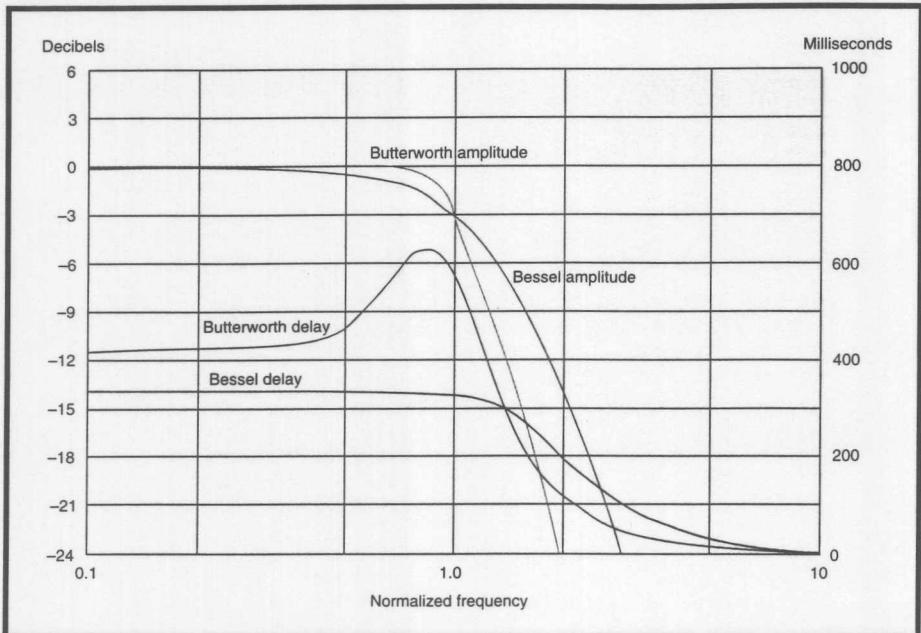


Figure 3—Four-pole Bessel and Butterworth filters are compared here. The Butterworth filter has a sharper corner and the Bessel filter has a constant delay.

ly. Frequencies below -3 dB are attenuated more than in a Butterworth filter. Frequencies above the cutoff frequency are attenuated less. Figure 3 contrasts these two filters.

You can use an Chebychev or elliptical filter if you want a sharp frequency cutoff. The elliptical filter achieves a sharper cutoff at the price of a low-frequency response that isn't flat. The Chebychev filter uses one or more notch sections to achieve a sharp cutoff. This can be useful when there is a particular frequency, such as a carrier, which you want to eliminate. Both of these filters have a dreadful effect on the shapes of pulses.

BUILDING FILTERS

Mathematically the behavior of a filter can be expressed by the ratio of two polynomials in s , the operator that represents differentiation with respect to time. The phase and amplitude response of the filter can be derived by replacing s with $j\omega$, where j is the square root of -1 and ω is 2π times the frequency in hertz.

In a low-pass filter, the numerator is usually unity. The highest power of s in the denominator governs the ultimate rate at which the filter response falls with frequency. The asymptotic slope is 6 dB per octave times the highest power of s .

For mathematical reasons, the highest power of s in the denominator is referred to as the number of "poles" of the filter and the highest power of s in the numerator is the number of its "zeros." These numbers make a handy way of classifying filters; for example, a filter with only poles is a low-pass filter. The more the number of poles exceeds the number of zeros, the steeper the ultimate filter response. Unfortunately, increasing complexity makes the filter performance critically dependent on the precise component values. Eight poles represent the practical limit for most purposes.

The constants multiplying the powers of s can be derived from the component values and, with somewhat

more difficulty, the component values can be calculated from the polynomials. Analysis programs such as the shareware program XFUNC simplify these conversions immensely. MicroCap IV also is helpful in analyzing the performance of filters.

NORMALIZATION

The idea behind normalization is that you don't have to design every filter from scratch. For any given configuration, you compute the component values once for some standard frequency and impedance. For example, you can select 1 rad/s or 1 Hz as the standard frequency and 1 Ω as the standard impedance. If you use 1 rad/s as the reference, the numerical values often turn out to be ratios of small integers. This is the standard I used, although using 1 Hz may be handier in practice. It is useful to normalize low-pass filters to unity DC gain (in other words, the numerical term in the denominator is unity).

A filter for any particular application can be calculated quickly from these standard values. First, the inductor and capacitor values are divided by the desired frequency (in radians per second). Next, the inductor and resistor values are multiplied by the desired impedance and the capacitor values are divided by the impedance. Of course, the values you get rarely correspond to conveniently available components, but often it's possible to juggle the impedance to make at least the capacitor values come out right. Tunable inductors are common enough components.

Two-pole Butterworth:	$\frac{1}{s^2 + 1.41421s + 1}$	C1 = 1.41421 C2 = 0.70711
Four-pole Butterworth:	$\frac{1}{(s^2 + 0.76537s + 1)(s^2 + 1.84776s + 1)}$	C1 = 2.61313 C2 = 0.38268 C3 = 1.08239 C4 = 0.92388
Two-pole Bessel:	$\frac{1}{0.61803s^2 + 1.36165s + 1}$	C1 = 0.68083 C2 = 0.90777
Four-pole Bessel:	$\frac{1}{(0.38899s^2 + 0.77425s + 1)(0.48890s^2 + 1.33966s + 1)}$	C1 = 1.00481 C2 = 0.38713 C3 = 0.72989 C4 = 0.66983

Figure 4—Equations and capacitors for active filters are normalized to -3 dB at 1 rad/s.

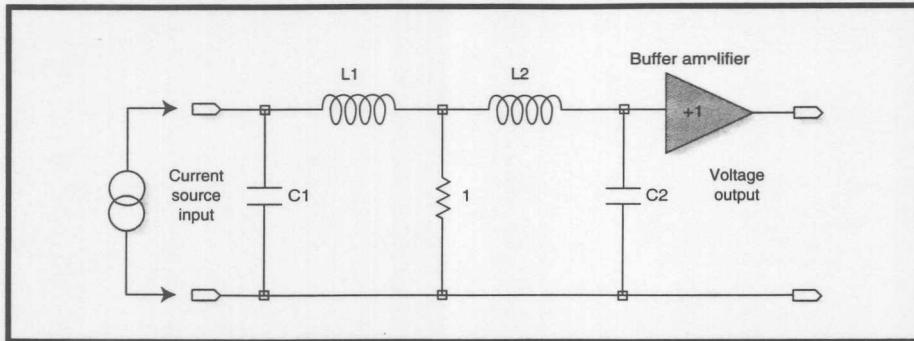


Figure 5—For fixed frequency applications, this four-pole LC filter is particularly easy to design. It must be driven from a current source and be followed by a voltage buffer.

DESIGNING FILTERS

Both the simple LCR filter and basic active filter have two poles. In most cases, a multi-pole filter can be synthesized from a series of two-pole sections. If you want a filter with an odd number of poles, you must add an extra RC roll-off. If the denominator of a filter can be factored into quadratic (two-pole) terms, the design becomes easy. Whereas LCR filters can't be readily stacked in series, active filters can be.

Figure 4 gives the equations and the normalized component values for two types of two- and four-pole active filters. I recalculated all the values from the first principles (published tables sometimes contain significant inaccuracies). Of course, component tolerances create errors of a few percent. These generally don't matter, but it's comforting to start from reliable inputs.

The equations are given for a frequency of 1 rad/s. Take this into account when calculating the component values for a particular frequency. Suppose you wanted a two-pole Butterworth filter with a -3-dB frequency of 10 kHz and decided that 10 kΩ would be a nice resistor value. (By the way, changing the 1:1 ratio of the resistors changes the damping factor, but usually isn't worth the bother.)

For a Butterworth filter, the damping factor (Z) is 0.7071,

$$\frac{1}{\sqrt{2}}$$

The normalized product of the capacitors is 1 and their ratio is 1:2. That is, if the resistors are 1 Ω and the corner frequency is 1 rad/s, then C1 is 0.7071F and C2 is 1.4142F.

INPUTS AND OUTPUTS

Most textbook filter designs assume that the input and output of the filter are voltages. Some configurations convert currents to voltages or vice versa. This is helpful when the related equipment supplies or requires a current. Fast DAC chips have current outputs and may need filtering to remove steps at the clock frequency. That's when a current-to-voltage filter is useful, particularly because the output capacitance of the DAC chip can be absorbed into the filter's input capacitance.

The same trick is handy when putting an anti-aliasing filter in front of a flash ADC. The ADC's input capacitance, which would otherwise be a significant load on the driving circuit, becomes an inherent part of the filter.

The classic two-pole passive LCR filter is a transimpedance device that converts a current to a voltage or vice versa. The current is treated as a voltage because it flows through the resistor.

Two LCR stages put back to back, sharing a common damping resistor, results in a practical four-pole filter

Actually, you want the corner frequency to equal 10 kHz, 62,832 rad/s. To calculate the real capacitor values, you divide the normalized values by 62,832 to get to 10 kHz and by another 10,000 to make the resistors 10 kΩ. These calculations make $C_1 = 1125 \text{ pF}$ and $C_2 = 2251 \text{ pF}$, both reasonable values. In practice, you would use, say, 2200-pF capacitors for C_2 and 1000-pF and 100-pF capacitors in parallel for C_1 . The resistors should be changed to 10.2 kΩ to compensate for the capacitors' values.

Be a Network HERO!



NetBurner's development kit
can put YOUR embedded
product on the Net.

Make Marketing say: **WOW!**



DevKit \$499

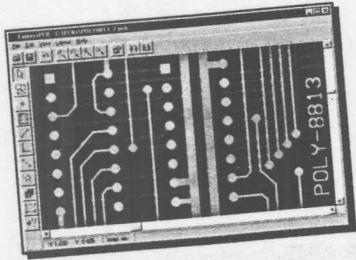
- Control and monitor anything
- Communicate through a local network, the Internet, or a telephone line
- Create a virtual front panel user interface and access it from anywhere in the world using a standard web browser
- OEM hardware from \$199

Buy a NetBurner today,
be networking tomorrow

Visit www.netburner.com
or Call: 858-530-0293

NetBurner
Networking in 1 Day!

- 1 Download our board layout software
- 2 Design your 2-sided plated-through PCB
- 3 Send us your layout over the Internet
- 4 In 2-3 business days, UPS delivers your boards, often under \$100



www.expresspcb.com

ExpressPCB announces \$59 Circuit Boards!

Select our MiniBoard service and get 3 top quality 2.5" x 3.8" PCBs for \$59 (free shipping)

www.expresspcb.com



PCBexpress
www.pcbexpress.com

- No tooling charge!
- Lot charges start at \$80
- Simple order process
- Quickturn, low quantities

TWO SERVICES FOR CIRCUIT BOARDS

ECD

INSTANT ON-LINE QUOTES!
(No sign-up required)

www.pcbexpress.com

- Quick Price Comparisons
- More options and added features
- Prototype & production quantities

For a discount on your next order enter the code: CC1810

amplifier to drive whatever follows it. Figure 5 shows this useful filter. You may recognize it as the one I used on the DAC output of my NCO generator project (*Circuit Cellar* 89, 90). It is easy to design because the resistor effectively isolates the two halves of the filter. Table 1 shows the normalized component values for this filter.

TUNING FILTERS

So far I've been assuming that the filter components never change. When equipment must operate over a range of input frequencies, you need a tunable filter. Luckily the Sallen and Key configuration is readily tuned by changing its two resistors in unison. The tuning can be electrically controlled over several decades. One way to do it is by using CdS photoresistors illuminated by LEDs as the tuning elements.

You do not need to use variable resistors. A switched-capacitor filter emulates the resistors by sending pulses of charge through capacitors. Changing the pulse rate changes the filter's cutoff frequency.

The most versatile tuning components are analog multiplier chips. [1] Many have a voltage input and current output, which are ideal for driving an integrator. Two integrators in the so-called "state variable" configuration make a versatile two-pole filter. It has simultaneous high-, band-, and low-pass outputs. These can be combined to make special-purpose filters such as a tunable notch filter. Changing the control voltage applied to the multipliers linearly tunes the cutoff frequency. Multiplier chips often have differential inputs that can simplify the filter. [2]

Any of these devices are capable of controlling the cutoff frequency of a filter over one or two decades. Wider ranges can be achieved by switching the tuning capacitors with relays or analog switches. I have designed equipment that incorporates pulse-shaping filters tunable over a six-decade range. [3] Traditionally, active filters were limited to cutoff frequencies less than 100 kHz. With modern

FILTERED OUT

I can't cram a filter design course into one article, but I hope I've given you an understanding of what types of analog filters exist and what they can do. The best sources for more information are trade magazines and application notes. With one exception, *The Art of Electronics*, there's a void in the textbook market. [5] Many books discuss the elementary LCR filter. Beyond those are the highly mathematical theoretical treatments, which overlook that someone might actually want to build a filter for a real application. One last thing I should tell you is that practical topics like filters that use nonideal auxiliary components or nonstandard configurations are omitted from both sources. ■

Tom Napier was a principal engineer in the Signal Recovery Group of Aydin Corp. He designed tunable pulse shaping filters into the single-board bit synchronizer selected by NASA for use in Space Shuttles. He is now a consultant and writer.

SOURCE

Micro-Cap IV

Rainbow Software

+44 0 208 2954500

www.micro-cap.co.uk

REFERENCES

- [1] C. Siu, "Design innovations provide for voltage-tunable, state-variable active filters for megahertz ranges," *EDN*, September 28, 1995.
- [2] T. Napier, "Multipliers implement tunable filters," *EDN*, March 16, 1992.
- [3] —, "Sifting Signals from Noise," *Electronics Now Magazine*, April 1999.
- [4] —, "Take tunable low-pass filters to new heights," *EDN*, January 15, 1998.
- [5] W. Hill and P. Horowitz, *The Art of Electronics*, 2nd ed., Cambridge University Press, Cambridge, England, 1989.